

MARKOV RANDOM FIELDS AND ITERATED TORIC FIBRE PRODUCTS

JAN DRAISMA AND FLORIAN M. OOSTERHOF

ABSTRACT. We prove that iterated toric fibre products from a finite collection of toric varieties are defined by binomials of uniformly bounded degree. This implies that Markov random fields built up from a finite collection of finite graphs have uniformly bounded Markov degree.

1. INTRODUCTION AND MAIN RESULTS

The notion of *toric fibre product* of two toric varieties goes back to [Sul07]. It is of relevance in algebraic statistics since it captures algebraically the Markov random field on a graph obtained by glueing two graphs along a common subgraph; see [RS16] and also below. In [Sul07, RS16, RS14] it is proved that under certain conditions, one can explicitly construct a Markov basis for the large Markov random field from bases for the components. For related results see [Shi12, EKS14, KR14].

However, these conditions are not always satisfied. Nevertheless, in [RS16, Conjecture 56] the hope was raised that when building larger graphs by glueing copies from a finite collection of graphs along a common subgraph, there might be a uniform upper bound on the Markov degree of the models thus constructed, independent of how many copies of each graph are used. A special case of this conjecture was proved in the same paper [RS16, Theorem 54]. We will prove the conjecture in general, and along the way we will link it to recent work [SS14] in *representation stability*. Indeed, an important point we would like to make, apart from proving said conjecture, is that algebraic statistics is a natural source of problems in *asymptotic algebra*, to which ideas from representation stability apply. Our main theorems are reminiscent of Sam's recent stabilisation theorems on equations and higher syzygies for secant varieties of Veronese embeddings [Sam15, Sam16].

Markov random fields. Let $G = (N, E)$ be a finite, undirected, simple graph and for each node $j \in N$ let X_j be a random variable taking values in the finite set $[d_j] := \{1, \dots, d_j\}$. A joint probability distribution on $(X_j)_{j \in N}$ is said to satisfy the *local Markov properties* imposed by the graph if every X_j is conditionally independent of all X_k with k at distance at least 2 from j , given X_l for all neighbours l of j .

On the other hand, a joint probability distribution f on the X_j is said to *factorise according to G* if for each maximal clique C of G and configuration $\alpha \in \prod_{j \in C} [d_j]$ of the random variables labelled by C there exists an interaction parameter θ_α^C such that for each configuration $\beta \in \prod_{j \in N} [d_j]$ of all random variables of G :

$$f(\beta) = \prod_{C \in \text{mcl}(G)} \theta_{\beta|_C}^C$$

where $\text{mcl}(G)$ is the set of maximal cliques of G , and $\beta|_C$ is the restriction of β to C .


These two notions are connected by the Hammersly-Clifford theorem, which says that a positive joint probability distribution on G factorises according to G if and only if it satisfies the Markov properties; see [HC71] or [Lau98, Theorem 3.9].

The set of all positive joint probability distributions on G that satisfy the Markov properties is therefore a subset of the image of the following map:

$$\varphi_G : \mathbb{C}^{\prod_{\text{Cemcl}(G)} \prod_{j \in C} [d_j]} \rightarrow \mathbb{C}^{\prod_{j \in N} [d_j]}, \quad (\theta_\alpha^C)_{C, \alpha} \mapsto \left(\prod_{\text{Cemcl}(G)} \theta_{\beta|C}^C \right)_\beta$$

It is the ideal I_G of polynomials vanishing on $\text{im } \varphi_G$ that is of interest in algebraic statistics. Since the components of φ_G are monomials, I_G is generated by finitely many binomials (differences of two monomials) in the standard coordinates on $\mathbb{C}^{\prod_{j \in N} [d_j]}$, and any finite generating set of binomials can be used to set up a Markov chain for testing whether given observations of the variables X_j are compatible with the assumption that their joint distribution factorises according to the graph G [DS98]. Both the ideal I_G and its zero locus are referred to as the *graphical model* of G .

Now suppose we have graphs G_1, \dots, G_s with node sets N_1, \dots, N_s , that $N_i \cap N_k$ equals a fixed set N_0 for all $i \neq k$ in $[s]$, and that the graph induced on N_0 by each G_i is equal to a fixed graph H . Moreover, for each $j \in \bigcup_i N_i$ fix a number d_j of states. We can then glue copies of the G_i along their common subgraph H , by which we mean first taking disjoint copies of the G_i and then identifying the nodes labelled by a fixed $j \in N_0$ across all copies. For $a_1, \dots, a_s \in \mathbb{Z}_{\geq 0}$, we denote the graph obtained by glueing a_i copies of graph $G_i, i \in [s]$ by $\sum_H^{a_1} G_1 +_H \dots +_H \sum_H^{a_s} G_s$:



Theorem 1. *Let G_1, \dots, G_s be graphs with a common subgraph H and a number of states associated to each node. Then there exists a uniform bound $D \in \mathbb{Z}_{\geq 0}$ such that for all multiplicities a_1, \dots, a_s , the ideal I_G of $G = \sum_H^{a_1} G_1 +_H \dots +_H \sum_H^{a_s} G_s$ is generated by binomials of degree at most D .*

Our proof will show that one needs only finitely many combinatorial types of binomials, independent of a_1, \dots, a_s , to generate I_G .

Example 2. In [RS14] it is proved that the ideal I_G for the complete bipartite graph $G = K_{3,N}$, with two states for each of the random variables, is generated in degree at most 12 for all N . The graph G is obtained by glueing N copies of $K_{3,1}$ along their common subgraph consisting of 3 nodes without any edges.

We will derive Theorem 1 from a general stabilisation result on *toric fibre products*, which we introduce next.

Toric fibre products. Fix a ground field K , let r be a natural number, and let $U_1, \dots, U_r, V_1, \dots, V_r$ be finite-dimensional vector spaces over K .

Define a bilinear operation

$$(1) \quad \prod_j U_j \times \prod_j V_j \rightarrow \prod_j (U_j \otimes V_j), \quad (u, v) \mapsto u * v := (u_1 \otimes v_1, \dots, u_r \otimes v_r).$$

Definition 3. The *toric fibre product* $X * Y$ of Zariski-closed subsets $X \subseteq \prod_j U_j$ and $Y \subseteq \prod_j V_j$ equals the Zariski-closure of the set $\{u * v \mid u \in X, v \in Y\}$. ♣

The product is associative and commutative up to reordering tensor factors. We can iterate this construction and form products like $X^{*2} * Y * Z^{*3}$, where Z also lives in a product of r vector spaces W_j . This variety lives in $\prod_j (U_j^{\otimes 2}) \otimes V_j \otimes (W_j^{\otimes 3})$.

We will not be taking toric fibre products of general varieties, but rather *Hadamard-stable* ones. For this, we have to choose coordinates on each U_j , so that $U_j = K^{d_j}$.

Definition 4. On K^d the Hadamard product is defined as $(a_1, \dots, a_d) \circ (b_1, \dots, b_d) = (a_1 b_1, \dots, a_d b_d)$. On $U := U_1 \times \dots \times U_r$, where $U_j = K^{d_j}$, it is defined component-wise. A set $X \subseteq \prod_j U_j$ is called *Hadamard-stable* if X contains the all-one vector $\mathbf{1}_U$ (the unit element of \circ) and if moreover $x \circ z \in X$ for all $x, z \in X$. ♣

In particular, the Zariski-closure in U of a subtorus of the $\sum_j d_j$ -dimensional torus $\prod_j (K \setminus \{0\})^{d_j}$ is Hadamard-stable. These are the toric varieties from the abstract.

Suppose that we also fix identifications $V_j = K^{d'_j}$ and a corresponding Hadamard multiplication $\prod_j V_j \times \prod_j V_j \rightarrow \prod_j V_j$. Equipping the spaces $U_j \otimes V_j$ with the natural coordinates and corresponding Hadamard multiplication, the two operations just defined satisfy $(u \circ u') * (v \circ v') = (u * v) \circ (u' * v')$ as well as $\mathbf{1}_{U_j \otimes V_j} = \mathbf{1}_{U_j} \otimes \mathbf{1}_{V_j}$. Consequently, if both $X \subseteq \prod_j U_j$ and $Y \subseteq \prod_j V_j$ are Zariski-closed and Hadamard-stable, then so is their toric fibre product $X * Y$. We can now formulate our second main result.

Theorem 5. Let $s, r \in \mathbb{Z}_{\geq 0}$. For each $i \in [s]$ and $j \in [r]$ let $d_{ij} \in \mathbb{Z}_{\geq 0}$ and set $V_{ij} := K^{d_{ij}}$. For each $i \in [s]$, let $X_i \subseteq \prod_j V_{ij}$ be a Hadamard-stable Zariski-closed subset. Then there exists a uniform bound $D \in \mathbb{Z}_{\geq 0}$ such that for any exponents a_1, \dots, a_s the ideal of $X_1^{*a_1} * \dots * X_s^{*a_s} \subseteq \prod_{j=1}^r \bigotimes_{i=1}^s V_{ij}^{\otimes a_i}$ is generated by polynomials of degree at most D .

Remark 6. A straightforward generalisation of this theorem also holds, where each X_i is a closed sub-scheme given by some ideal J_i in the coordinate ring of $\prod_j V_{ij}$. The notion of Hadamard-stable then says that the pull-back of J_i lies in the ideal of $X_i \times X_i$, and the toric fibre product is defined in [Sul07, RS16]. But since this generality would slightly obscure our arguments, we have decided to present explicitly the version with Zariski-closed subsets.

Also, the theorem remains valid if we remove the condition that the X_i contain the all-one vector, but require only that they be closed under Hadamard-multiplication; see Remark 19. ♣

Organisation of this paper. The remainder of this paper is organised as follows. In Section 2 we introduce the categories of (affine) **Fin**-varieties and, dually, **Fin**^{op}-algebras. The point is that, as we will see in Section 4, the iterated toric fibre products together form such a **Fin**-variety (or rather a **Fin**^s-variety, where **Fin**^s is the product category of s copies of **Fin**).

Indeed, they sit naturally in a Cartesian product of copies of the **Fin**-variety of rank-one tensors, which, as we prove in Section 3, is Noetherian (Theorem 10). This Noetherianity result is of a similar flavour as the recent result from [SS14] (see also [DK14, Proposition 7.5] which follows the same proof strategy) that any finitely generated **Fin**^{op}-module is Noetherian; this result played a crucial role in a proof of the *Artinian conjecture*. However, our Noetherianity result concerns certain **Fin**^{op}-algebras rather than modules, and is more complicated. Finally, in Section 4 we first prove Theorem 5 and then derive Theorem 1 as a corollary.

Acknowledgements. Both authors are (partially) supported by Draisma's *Vici* grant *Stabilisation in Algebra and Geometry* from the Netherlands Organisation for Scientific Research. The current paper is partly based on the second author's Master's thesis at Eindhoven University of Technology; see [Oos16]. We thank Johannes Rauh and Seth Sullivant for fruitful discussions at *Prague Stochastics*, August 2014.

2. AFFINE **Fin**-VARIETIES AND **Fin**^{op}-ALGEBRAS

The category **Fin** has as objects all finite sets and as morphisms all maps. Its opposite category is denoted **Fin**^{op}. When **C** is another category whose objects are called *somethings*, then a **Fin**-*something* is a covariant functor from **Fin** to **C** and a **Fin**^{op}-*something* is a contravariant functor from **Fin** to **C**. The **C**-homomorphism associated to a **Fin**-morphism π is denoted π_* (the *push-forward* of π) in the covariant case and π^* (the *pull-back* of π) in the contravariant case.

More generally, we can replace **Fin** by any category **D**. The **D**-somethings themselves form a category, in which morphisms are natural transformations. In our paper, **D** will be closely related to **Fin** or to **Fin**^s, the *s*-fold product of **Fin**.

Here are three instances of **Fin** or **Fin**^{op}-somethings crucial to our paper.

Example 7. Fix $n \in \mathbb{Z}_{\geq 0}$. Then the functor $S \mapsto [n]^S$ is a **Fin**^{op}-set, which to $\pi \in \text{Hom}_{\mathbf{Fin}}(S, T)$ associates the map $\pi^* : [n]^T \rightarrow [n]^S, \alpha \mapsto \alpha\pi$, the composition of α and π .

Building on this example, the functor $A_n : S \mapsto K[x_\alpha \mid \alpha \in [n]^S]$, a polynomial ring in variables labelled by $[n]^S$, is a **Fin**^{op}-*K*-algebra, which associates to π the *K*-algebra homomorphism $\pi^* : A(T) \rightarrow A(S), x_\alpha \mapsto x_{\alpha\pi}$.

Third, we define an affine **Fin**-*K*-vector space Q_n by $Q_n : S \mapsto (K^n)^{\otimes S}$, the space of $n \times \cdots \times n$ -tensors with factors labelled by *S*, which sends $\pi \in \text{Hom}_{\mathbf{Fin}}(S, T)$ to the linear morphism $\pi_* : (K^n)^{\otimes S} \rightarrow (K^n)^{\otimes T}$ determined by

$$\pi_* : \otimes_{i \in S} v_i \mapsto \otimes_{j \in T} \left(\bigcirc_{i \in \pi^{-1}(j)} v_i \right),$$

where \bigcirc is the Hadamard product in K^n . We follow the natural convention that an empty Hadamard product equals the all-one vector $\mathbf{1} \in K^n$; in particular, this holds in the previous formula for all $j \in T$ that are not in the image of π .

The ring A_n and the space Q_n are related as follows: $Q_n(S)$ has a basis consisting of vectors $e_\alpha := \otimes_{i \in S} e_{\alpha_i}, \alpha \in [n]^S$, where e_1, \dots, e_n is the standard basis of K^n ; $A_n(S)$ is the coordinate ring of $Q_n(S)$ generated by the dual basis $(x_\alpha)_{\alpha \in [n]^S}$; and for $\pi \in \text{Hom}_{\mathbf{Fin}}(S, T)$ the pullback $\pi^* : A_n(T) \rightarrow A_n(S)$ is the homomorphism of

K -algebras dual to the linear map $\pi_* : Q_n(S) \rightarrow Q_n(T)$. Indeed, this is verified by the following computation for $\alpha \in [n]^T$:

$$x_\alpha(\pi_* \otimes_{i \in S} v_i) = x_\alpha(\otimes_{j \in T} (\bigcirc_{i \in \pi^{-1}(j)} v_i)) = \prod_{j \in T} (\bigcirc_{i \in \pi^{-1}(j)} v_i)_{\alpha_j} = \prod_{i \in S} (v_i)_{\alpha_{\pi(i)}} = \pi^*(x_\alpha)(\otimes_{i \in S} v_i).$$

This is used in Section 3. ♣

In general, by *algebra* we shall mean an associative, commutative K -algebra with 1, and homomorphisms are required to preserve 1. So a **Fin**^{op}-algebra B assigns to each finite set S an algebra and to each map $\pi : S \rightarrow T$ an algebra homomorphism $\pi^* : B(T) \rightarrow B(S)$. An *ideal* in B is a **Fin**^{op}-subset I of B (i.e., $I(S)$ is a subset of $B(S)$ for each finite set S and π^* maps $I(T)$ into $I(S)$) such that each $I(S)$ is an ideal in $B(S)$; then $S \mapsto B(S)/I(S)$ is again a **Fin**^{op}-algebra, the *quotient* B/I of B by I .

Given a **Fin**^{op}-algebra B , finite sets S_j for j in some index set J , and an element $b_j \in B(S_j)$ for each j , there is a unique smallest ideal I in B such that each $I(S_j)$ contains b_j . This ideal is constructed as:

$$I(S) = (\pi^*(b_j) \mid j \in J, \pi \in \text{Hom}(S, S_j))$$

This is the ideal *generated* by the b_j . A **Fin**^{op}-algebra is called *Noetherian* if each ideal I in it is generated by finitely many elements in various $I(S_j)$, i.e., J can be taken finite.

Example 8. The **Fin**^{op}-algebra A_1 is Noetherian. Indeed, $A_1(S)$ is the polynomial ring $K[t]$ in a single variable t for all S , and the homomorphism $A_1(T) \rightarrow A_1(S)$ is the identity $K[t] \rightarrow K[t]$. So Noetherianity follows from Noetherianity of the algebra $K[t]$.

For $n \geq 2$ the **Fin**^{op}-algebra A_n is *not* Noetherian. For instance, consider the monomials

$$u_2 := x_{21}x_{12} \in A_n([2]), u_3 := x_{211}x_{121}x_{112} \in A_n([3]), u_4 := x_{2111}x_{1211}x_{1121}x_{1112} \in A_n([4]),$$

and so on. For any map $\pi : [k] \rightarrow [l]$ with $k > l$, by the pigeon hole principle there exist two indices $i, j \in [k]$ such that $\pi(i) = \pi(j) =: m \in [l]$. Then $\pi^*x_{1\dots 121\dots 1}$, where the 2 is in the m -th position, is a variable with at least two indices equal to 2. Since u_k contains no such variable, π^*u_l does not divide u_k . So u_2, u_3, \dots generates a non-finitely generated monomial **Fin**^{op}-ideal in A_n . (On the other hand, for each d the piece of A_n of homogeneous polynomials of degree at most d is Noetherian as a **Fin**^{op}-module, see [DK14, Proposition 7.5].) ♣

We shall see in the following section that certain interesting quotients of each A_n are Noetherian.

3. RANK-ONE TENSORS FORM A NOETHERIAN **Fin**-VARIETY

Recall that the **Fin**^{op}-algebra A_n assigns to S the coordinate ring of the tensor product $Q_n(S) = (K^n)^{\otimes S}$. Let $Q_n^{\leq 1}(S)$ be the variety of *rank-one tensors*, i.e., those of the form $\otimes_{i \in S} v_i$ for vectors $v_i \in K^n$. We claim that this defines a Zariski-closed **Fin**-subvariety $Q_n^{\leq 1}$ of Q_n .

For this, we must verify that for a map $\pi : S \rightarrow T$ the map $Q_n(S) \rightarrow Q_n(T)$ dual to the algebra homomorphism $A_n(T) \rightarrow A_n(S)$ sends $Q_n^{\leq 1}(S)$ into $Q_n^{\leq 1}(T)$. And indeed, in Example 7 we have seen that this map sends

$$\otimes_{i \in S} v_i \mapsto \otimes_{j \in T} \left(\bigcirc_{i \in \pi^{-1}(j)} v_i \right),$$

where \bigcirc is the Hadamard product.

It is well known that (if K is infinite) the ideal in $A_n(S)$ of $Q_n^{\leq 1}(S)$ equals the ideal $I_n(S)$ generated by all binomials constructed as follows. Partition S into two parts S_1, S_2 , let $\alpha_i, \beta_i \in [n]^{S_i}$ and write $\alpha_1 \parallel \alpha_2$ for the element of $[n]^S$ which equals α_i on S_i . Then we have the binomial

$$(2) \quad x_{\alpha_1 \parallel \alpha_2} x_{\beta_1 \parallel \beta_2} - x_{\alpha_1 \parallel \beta_2} x_{\beta_1 \parallel \alpha_2} \in I_n(S),$$

and $I_n(S)$ is the ideal generated by these for all partitions and all $\alpha_1, \alpha_2, \beta_1, \beta_2$. The functor $S \mapsto I_n(S)$ is an ideal in the \mathbf{Fin}^{op} -algebra A_n ; for infinite K this follows from the computation above, and for arbitrary K it follows since the binomials above are mapped to binomials by pull-backs of maps $T \rightarrow S$ in \mathbf{Fin} . Moreover, I_n is finitely generated (see also [DK14, Lemma 7.4]):

Lemma 9. *The ideal I_n in the \mathbf{Fin}^{op} -algebra A_n is finitely generated.*

Proof. In the determinantal equation (2), if there exist distinct $j, l \in S_1$ such that $\alpha_1(j) = \alpha_2(l)$ and $\beta_1(j) = \beta_2(l)$, then the equation comes from an equation in $I_n(S \setminus \{j\})$ via the map $S \rightarrow S \setminus \{j\}$ that is the identity on $S \setminus \{j\}$ and maps $j \rightarrow l$. By the pigeon hole principle this happens when $|S_1| > n^2$. Similarly for $|S_2| > n^2$. Hence I_n is certainly generated by $I_n([2n^2 - 1])$. \square

The main result in this section is the following.

Theorem 10. *For each $n \in \mathbb{Z}_{\geq 0}$ the coordinate ring A_n/I_n of the \mathbf{Fin} -variety $Q_n^{\leq 1}$ of rank-one tensors is a Noetherian \mathbf{Fin}^{op} -algebra.*

Our proof follows the general technique from [SS14], namely, to pass to a suitable category close to \mathbf{Fin} that allows for a Gröbner basis argument. However, the relevant well-partial-orderedness proved below is new and quite subtle. We use the category \mathbf{OS} from [SS14] (also implicit in [DK14, Section 7]) defined as follows.

Definition 11. The objects of the category \mathbf{OS} (“ordered-surjective”) are all finite sets equipped with a linear order and the morphisms $\pi : S \rightarrow T$ are all surjective maps with the additional property that the function $T \rightarrow S$, $j \mapsto \min \pi^{-1}(j)$ is strictly increasing. \clubsuit

Any \mathbf{Fin} -algebra is also an \mathbf{OS} -algebra, and \mathbf{OS} -Noetherianity implies \mathbf{Fin} -Noetherianity. So to prove Theorem 10 we set out to prove the stronger statement that A_n/I_n is, in fact, \mathbf{OS} -Noetherian.

We get a more concrete grip on the K -algebra A_n/I_n through the following construction. Let M_n denote the (Abelian) \mathbf{Fin}^{op} -monoid defined by

$$M_n(S) := \left\{ \alpha \in \mathbb{Z}_{\geq 0}^{[n] \times S} \mid \forall j, l \in S : \sum_{i=1}^n \alpha_{ij} = \sum_{i=1}^n \alpha_{il} \right\},$$

in which the multiplication is given by addition, and where the pull-back of a map $\pi : S \rightarrow T$ is the map $\pi^* : M_n(T) \rightarrow M_n(S)$ sending $(\alpha_{ij})_{i \in [n], j \in T}$ to $(\alpha_{i\pi(j)})_{i \in [n], j \in S}$. So elements of $M_n(S)$ are matrices with nonnegative integral entries and constant column sum. Let KM_n denote the **Fin**-algebra sending S to the monoid K -algebra $KM_n(S)$. The following proposition is a reformulation of a well-known fact.

Proposition 12. *The \mathbf{Fin}^{op} -algebra A_n/I_n is isomorphic to the \mathbf{Fin}^{op} -algebra KM_n (and the same is true when both are regarded as \mathbf{OS}^{op} -algebras).*

Proof. For each finite set S , the K -algebra homomorphism $\Phi_S : A_n(S) \rightarrow KM_n(S)$ that sends $x_\alpha, \alpha \in [n]^S$ to the $[n] \times S$ matrix in $M_n(S)$ with a 1 at the positions (α_j, j) , $j \in S$ and zeroes elsewhere is surjective and has as kernel the ideal $I_n(S)$. Moreover, if $\pi : S \rightarrow T$ is a morphism in **Fin**, then we have $\Phi_T \pi^* = \pi^* \Phi_S$, i.e., the $(\Phi_S)_S$ define a natural transformation. \square

Choose any monomial order $>$ on $\mathbb{Z}_{\geq 0}^n$, i.e., a well-order such that $a > b$ implies $a + c > b + c$ for every $a, b, c \in \mathbb{Z}_{\geq 0}^n$. Then for each object S in **OS** we define a linear order $>$ on $M_n(S)$, as follows: $\alpha > \beta$ if $\alpha \neq \beta$ and the smallest $j \in S$ with $\alpha_{.,j} \neq \beta_{.,j}$ (i.e., the j -th column of α is not equal to that of β) satisfies $\alpha_{.,j} > \beta_{.,j}$ in the chosen monomial order on $\mathbb{Z}_{\geq 0}^n$. A straightforward verification shows that this is a monomial order on $M_n(S)$ (we will call the elements of $M_n(S)$ monomials, even though $KM_n(S)$ is not a polynomial ring). Moreover, for various S , these orders are interrelated as follows.

Lemma 13. *For any $\pi \in \text{Hom}_{\mathbf{OS}}(S, T)$ and $\alpha, \beta \in M_n(T)$, we have $\alpha > \beta \Rightarrow \pi^* \alpha > \pi^* \beta$.*

Proof. If $j \in T$ is the smallest column index where α and β differ, then $\alpha' := \pi^* \alpha$ and $\beta' := \pi^* \beta$ differ in column $l := \min \pi^{-1}(j)$, where they equal $\alpha_{.,j}$ and $\beta_{.,j}$, respectively, and the former is larger than the latter. Furthermore, if l' is the smallest position where α', β' differ, then $\alpha_{.,\pi(l')} \neq \beta_{.,\pi(l')}$ and hence $\pi(l') \geq j$ and hence $l' = \min \pi^{-1}(\pi(l')) \geq \min \pi^{-1}(j) = l$. Hence in fact $l = l'$ and $\pi^* \alpha > \pi^* \beta$. \square

In addition to the well-order \leq on each individual $M_n(S)$, we also need the following partial order $|$ on the union of all of them.

Definition 14. Let S, T be objects in **OS**. We say that $\alpha \in M_n(T)$ *divides* $\beta \in M_n(S)$ if there exist a $\pi \in \text{Hom}_{\mathbf{OS}}(S, T)$ and a $\gamma \in M_n(S)$ such that $\beta = \gamma + \pi^* \alpha$. In this case, we write $\alpha | \beta$. \clubsuit

The key combinatorial property of the relation just defined is the following.

Proposition 15. *The relation $|$ is a well-quasi-order, that is, for any sequence $\alpha^{(1)} \in M_n(S_1), \alpha^{(2)} \in M_n(S_2), \dots$ there exist $i < j$ such that $\alpha^{(i)} | \alpha^{(j)}$.*

Proof. First, to each $\alpha \in M_n(S)$ we associate the monomial ideal $P(\alpha)$ in the polynomial ring $R := K[z_1, \dots, z_n]$ (here K is but a place holder) generated by the monomials $z^{\alpha_{.,j}}, j \in S$. The crucial fact that we will use about monomial ideals in R is that in any sequence P_1, P_2, \dots of such ideals there exist $i < j$ such that $P_i \supseteq P_j$ —in other words, monomial ideals are well-quasi-ordered with respect to reverse inclusion [Mac01].

To prove the proposition, suppose, on the contrary, that there exists a sequence as above with $\alpha^{(i)} \not\ll \alpha^{(j)}$ for all $i < j$. Such a sequence is called *bad*. Then by basic properties of well-quasi-orders, some bad sequence exists in which moreover

$$(3) \quad P(\alpha^{(1)}) \supseteq P(\alpha^{(2)}) \supseteq \dots$$

Among all bad sequences *with this additional property* choose one in which, for each $j = 1, 2, \dots$, the cardinality $|S_j|$ is minimal among all bad sequences starting with $\alpha^{(1)}, \dots, \alpha^{(j-1)}$.

Write $\alpha^{(j)} = (\gamma^{(j)} | \beta^{(j)})$, where $\beta^{(j)} \in \mathbb{Z}_{\geq 0}^n$ is the last column (the one labelled by the largest element of S_j), and $\gamma^{(j)}$ is the remainder. By Dickson's lemma, there exists a subsequence $j_1 < j_2 < \dots$ such that $\beta^{(j_1)}, \beta^{(j_2)}, \dots$ increase weakly in the coordinate-wise ordering on $\mathbb{Z}_{\geq 0}^n$. By restricting to a further subsequence, we may moreover assume that also

$$(4) \quad P(\gamma^{(j_1)}) \supseteq P(\gamma^{(j_2)}) \supseteq \dots$$

Then consider the new sequence

$$\alpha^{(1)}, \dots, \alpha^{(j_1-1)}, \gamma^{(j_1)}, \gamma^{(j_2)}, \dots$$

By (4) and (3), and since $P(\alpha^{(j)}) \supseteq P(\gamma^{(j)})$, this sequence also satisfies (3). We claim that, furthermore, it is bad.

Suppose, for instance, that $\gamma^{(j_1)} | \gamma^{(j_2)}$. Set $a_i := \max S_{j_i}$ for $i = 1, 2$. Then there exists a $\pi \in \text{Hom}_{\mathbf{OS}}(S_{j_2} \setminus \{a_2\}, S_{j_1} \setminus \{a_1\})$ such that $\gamma^{(j_2)} - \pi^* \gamma^{(j_1)} \in M_n(S_{j_2} \setminus \{a_2\})$. But then extend π to an element π of $\text{Hom}_{\mathbf{OS}}(S_{j_2}, S_{j_1})$ by setting $\pi(a_2) := a_1$; since $\beta^{(j_1)}$ is coordinate-wise smaller than $\beta^{(j_2)}$ we find that $\alpha^{(j_2)} - \pi^* \alpha^{(j_1)} \in M_n(S_{j_2})$, so $\alpha^{(j_1)} | \alpha^{(j_2)}$, a contradiction with the badness of the original sequence.

On the other hand, suppose for instance that $\alpha^{(1)} | \gamma^{(j_2)}$ and write $a_2 := \max S_{j_2}$. Then there exists a $\pi \in \text{Hom}_{\mathbf{OS}}(S_{j_2} \setminus \{a_2\}, S_1)$ such that $\gamma^{(j_2)} - \pi^* \alpha^{(1)} \in M_n(S_{j_2} \setminus \{a_2\})$. Now, *and this is why we required that (3) holds*, since $P(\alpha^{(j_2)}) \subseteq P(\alpha^{(1)})$, there exists an element $s \in S_1$ such that the column $\beta^{(j_2)}$ is coordinatewise at least as large as the s -th column of $\alpha^{(1)}$. Extend π to an element of $\text{Hom}_{\mathbf{OS}}(S_{j_2}, S_1)$ by setting $\pi(a_2) = s$. Since a_2 is the maximal element of S_{j_2} , this does not destroy the property that the function $\min \pi^{-1}(\cdot)$ be increasing in its argument. Moreover, this π has the property that $\alpha^{(j_2)} - \pi^* \alpha^{(1)} \in M_n(S_{j_2})$, again a contradiction.

Since we have found a bad sequence satisfying (3) but with strictly smaller underlying set at the j_1 -st position, we have arrived at a contradiction. \square

Next we follow a Gröbner basis argument.

Proof of Theorem 10. We prove the stronger statement that KM_n is Noetherian as an \mathbf{OS} -algebra. Let P be any ideal in $A_n/I_n = KM_n$. For each object S in \mathbf{OS} , let $L(S) \subseteq M_n(S)$ denote the set of leading terms of nonzero elements of $P(S)$ relative to the ordering $>$. Proposition 15 implies that there exists a finite collection S_1, \dots, S_N and $\alpha^{(j)} \in L(S_j)$ such that each element of each $L(S)$ is divisible by some $\alpha^{(j)}$. Correspondingly, there exist elements $f_j \in P(S_j)$ with leading monomial $\alpha^{(j)}$ and leading coefficient 1. To see that the f_j generate P , suppose that there exists an S such that $P(S)$ is not contained in the ideal generated by the f_j , and let $g \in P(S)$ have minimal leading term β among all elements of $P(S)$ not in the ideal generated

by the f_j ; without loss of generality g has leading coefficient 1. By construction, there exists some j and some $\pi \in \text{Hom}_{\text{OS}}(S, S_j)$ such that $\beta - \pi^* \alpha^{(j)} \in M_n(S)$. But now, by Lemma 13, we find that the leading monomial of $\pi^* f_j$ equals $\pi^* \alpha^{(j)}$, hence subtracting a monomial times $\pi^* f_j$ from g we obtain an element of $P(S)$ with smaller leading monomial that is not in the ideal generated by the f_j —a contradiction. \square

Below, we will need the following generalisation of Theorem 10.

Theorem 16. *For any $n_1, \dots, n_r \in (\mathbb{Z})_{\geq 0}$ the **Fin**-algebra (or **OS**-algebra) $(A_{n_1}/I_{n_1}) \otimes \dots \otimes (A_{n_r}/I_{n_r})$ is Noetherian.*

Proof. This algebra is isomorphic to $B := K(M_{n_1} \times \dots \times M_{n_r})$. There is a natural embedding $\iota : M_{n_1} \times \dots \times M_{n_r} \rightarrow M_{n_1 + \dots + n_r} =: M_n$ by forming a block matrix; its image consists of block matrices with constant *partial* column sums. And while a subalgebra of a Noetherian algebra is not necessarily Noetherian, this is true in the current setting.

The crucial point is that if $\alpha_i \in (M_{n_1} \times \dots \times M_{n_r})(S_i)$ for $i = 1, 2$, then *a priori* $\iota(\alpha_1) | \iota(\alpha_2)$ only means that $\iota(\alpha_2) - \pi^* \iota(\alpha_1) \in M_n(S_2)$; but since both summands have constant partial column sums, so does their difference, so in fact, the difference lies in the image of ι . With this observation, the proof above for the case where $r = 1$ goes through unaltered for arbitrary r . \square

Remark 17. Similar arguments for passing to sub-algebras are also used in [HS12] and [Dra10]. \clubsuit

4. PROOFS OF THE MAIN RESULTS

In this section we prove theorems 1 and 5.

Toric fibre products. To prove Theorem 5, we work with a product of s copies of the category **Fin**; one for each of the varieties X_i whose iterated fibre products are under consideration. Let $s, r \in \mathbb{Z}_{\geq 0}$. For each $i \in [s]$ and $j \in [r]$ let $d_{ij} \in \mathbb{Z}_{\geq 0}$ and set $V_{ij} := K^{d_{ij}}$. Consider the **Fin** ^{s} -variety \mathcal{V} that assigns to an s -tuple $S = (S_1, \dots, S_s)$ the product

$$\prod_{j=1}^r \bigotimes_{i=1}^s V_{ij}^{\otimes S_i}$$

and to a morphism $\pi = (\pi_1, \dots, \pi_s) : S \rightarrow T := (T_1, \dots, T_s)$ in **Fin** ^{s} the linear map $\mathcal{V}(S) \rightarrow \mathcal{V}(T)$ determined by

$$(5) \quad \left(\otimes_i \otimes_{k \in S_i} v_{ijk} \right)_{j \in [r]} \mapsto \left(\otimes_i \otimes_{l \in T_i} \left(\bigcirc_{k \in \pi_i^{-1}(l)} v_{ijk} \right) \right)_{j \in [r]},$$

where \bigcirc stands for the Hadamard product in V_{ij} . Let $\mathcal{Q}^{\leq 1}(S)$ be the Zariski-closed subset of $\mathcal{V}(S)$ consisting of r -tuples of tensors of rank at most one; thus $\mathcal{Q}^{\leq 1}$ is a **Fin** ^{s} -subvariety of \mathcal{V} .

For each $i \in [s]$ let $X_i \subseteq \prod_{j=1}^r V_{ij}$ be a Hadamard-stable Zariski-closed subset. Then for any tuple $S = (S_1, \dots, S_s)$ in **Fin** ^{s} the variety $\mathcal{X}(S) := X_1^{*S_1} * \dots * X_s^{*S_s}$ is a Zariski-closed subset of $\mathcal{V}(S)$.

Lemma 18. *The association $S \mapsto X(S)$ defines a \mathbf{Fin}^s -closed subvariety of $Q^{\leq 1}$.*

Proof. From the definition of $*$ in (1) it is clear that the elements of $X(S)$ are r -tuples of tensors of rank at most 1. Furthermore, for a morphism $S \rightarrow T$ in \mathbf{Fin}^s the linear map $\mathcal{V}(S) \rightarrow \mathcal{V}(T)$ from (5) sends $X(S)$ into $X(T)$ —here we use that if $(v_{i,j,k})_{j \in [r]} \in X_i$ for each $k \in \pi^{-1}(l)$, then also $(\bigcirc_{k \in \pi^{-1}(l)} v_{ijk})_{j \in [r]} \in X_i$ since X_i is Hadamard-stable. \square

Now Theorem 5 follows once we know that $Q^{\leq 1}$ is a Noetherian $(\mathbf{Fin}^s)^{\text{op}}$ -algebra. For $s = 1$ and $r = 1$ this is Theorem 10 with n equal to d_{11} . For $s = 1$ and general r , this is Theorem 16 with n_j equal to d_{1j} .

For $r = 1$ and general s , Theorem 5 follows from a \mathbf{Fin}^s -analogue of Theorem 10, which is proved as follows. The coordinate ring of $Q^{\leq 1}(S_1, \dots, S_s)$ is the subring of $K(M_{d_{11}}(S_1) \times \dots \times M_{d_{s1}}(S_s))$ spanned by the monomials corresponding to s -tuples of matrices with the *same* constant column sum. Using Proposition 15 and the fact that a finite product of well-quasi-ordered sets is well-quasi-ordered one finds that the natural \mathbf{Fin}^s -analogue on $M_{d_{11}} \times \dots \times M_{d_{s1}}$ of the division relation $|$ is a well-quasi-order; and this implies, once again, that $Q^{\leq 1}$ is a Noetherian $(\mathbf{OS}^s)^{\text{op}}$ -algebra.

Finally, for general s and general r , the result follows as in the proof of Theorem 16. This proves the Theorem 5 in full generality. \square

Remark 19. The only place where we used that the X_i contain the all-one vector is in the proof of Lemma 18 when $\pi^{-1}(l)$ happens to be empty. If we do not require this, then the conclusion of Theorem 5 still holds, since one can work directly with the category \mathbf{OS}^s in which morphisms π are surjective.

Markov random fields. Given a finite (undirected, simple) graph $G = (N, E)$ with a number d_j of states attached to each node $j \in N$, the graphical model is $X_G := \overline{\text{im } \varphi_G}$, where φ_G is the parameterisation

$$\varphi_G : \mathbb{C}^{\prod_{C \in \text{mcl}(G)} \prod_{j \in C} [d_j]} \rightarrow \mathbb{C}^{\prod_{j \in N} [d_j]}, \quad (\theta_\alpha^C)_{C, \alpha} \mapsto \left(\prod_{C \in \text{mcl}(G)} \theta_{\beta|_C}^C \right)_\beta.$$

Lemma 20. *For any finite graph G , the graphical model X_G is Hadamard-closed.*

Proof. The parameterisation φ sends the all-one vector in the domain to the all-one vector $\mathbf{1}$ in the target space, so $\mathbf{1} \in \text{im } \varphi$. Moreover, if θ, θ' are two parameter vectors, then $\varphi(\theta \circ \theta') = \varphi(\theta) \circ \varphi(\theta')$, so $\text{im } \varphi$ is Hadamard-closed. Then so is its closure. \square

Following [Sul07], we relate graph glueing to toric fibre products. We are given finite graphs G_1, \dots, G_s with node sets N_1, \dots, N_s such that $N_i \cap N_k = N_0$ for all $i \neq k$ in $[s]$ and such that each G_i induces the same graph H on N_0 . Moreover, for each $j \in \bigcup_i N_i$ we fix a number d_j of states.

For each $\beta_0 \in \prod_{j \in N_0} [d_j]$ and each $i \in [s]$ set $V_{i, \beta_0} := \mathbb{C}^{\prod_{j \in N_i \setminus N_0} [d_j]}$, which we interpret as the ambient space of the part of the probability table of the variables X_j , $j \in N_i$ where we have fixed the states of the variables in N_0 to β_0 —up to scaling, these are

the conditional joint probabilities for the $X_j, j \in N_i \setminus N_0$ given that the $X_j, j \in N_0$ are in joint state β_0 . For $\beta \in \prod_{j \in N_i} [d_j]$ write $\beta = \beta_0 \| \beta'$ where β_0, β' are the restrictions of β to N_0 and $N_i \setminus N_0$, respectively. For each maximal clique C in G_i define $C_0 = C \cap N_0$ and $C' = C \setminus N_0$. Correspondingly, decompose $\alpha \in \prod_{j \in C} [d_j]$ as $\alpha = \alpha_0 \| \alpha'$, where α_0, α' are the restrictions of α to C_0 and C' , respectively.

Then the graphical model $X_{G_i} := \overline{\text{im } \varphi_{G_i}}$ is the closure of the image of the parametrisation

$$\varphi_i : \mathbb{C}^{\prod_{C \in \text{mcl}(G_i)} \prod_{j \in C} [d_j]} \rightarrow \prod_{\beta_0 \in \prod_{j \in N_0} [d_j]} V_{i, \beta_0}, \quad (\theta_\alpha^C)_{C, \alpha} \mapsto \left(\left(\prod_{C \in \text{mcl}(G)} \theta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta'} \right)_{\beta_0}.$$

Setting $r := \prod_{j \in N_0} d_j$, we are exactly in the setting of the previous sections: for each $i, k \in [s]$, we have the bilinear map

$$* : \prod_{\beta_0} V_{i, \beta_0} \times \prod_{\beta_0} V_{k, \beta_0} \rightarrow \prod_{\beta_0} (V_{i, \beta_0} \otimes V_{k, \beta_0}), \quad ((v_{i, \beta_0})_{\beta_0}, (v_{k, \beta_0})_{\beta_0}) \mapsto (v_{i, \beta_0} \otimes v_{k, \beta_0})_{\beta_0},$$

and we can take iterated products of this type. The space on the right is naturally isomorphic to $\mathbb{C}^{\prod_{j \in N_1 \cup N_2} [d_j]}$, the space of probability tables for the joint distribution of the variables labelled by the vertices in the glued graph $G_i +_H G_k$. Under this identification we have the following.

Proposition 21. *For $G := \sum_H^{a_1} G_1 +_H \dots +_H \sum_H^{a_s} G_s$ we have $X_G = X_{G_1}^{*a_1} \dots X_{G_s}^{*a_s}$.*

Proof. It suffices to prove this for the gluing of two graphs. Note that a clique in $G := G_1 +_H G_2$ is contained entirely in either G_1 or G_2 , or in both but then already in H . Let θ, η be a parameter vectors in the domains of $\varphi_{G_1}, \varphi_{G_2}$, respectively. Then

$$\begin{aligned} & \varphi_{G_1}(\theta) * \varphi_{G_2}(\eta) \\ &= \left(\left(\prod_{C \in \text{mcl}(G_1)} \theta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta'} \right)_{\beta_0} * \left(\left(\prod_{C \in \text{mcl}(G_2)} \eta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta'} \right)_{\beta_0} \\ &= \left(\left(\prod_{C \in \text{mcl}(G_1)} \theta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta' \in \prod_{j \in N_1 \setminus N_0} [d_j]} \otimes \left(\prod_{C \in \text{mcl}(G_2)} \eta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta' \in \prod_{j \in N_2 \setminus N_0} [d_j]} \right)_{\beta_0} \\ &= \left(\left(\prod_{C \in \text{mcl}(G_1)} \theta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \cdot \prod_{C \in \text{mcl}(G_2)} \eta_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta' \in \prod_{j \in (N_1 \cup N_2) \setminus N_0} [d_j]} \right)_{\beta_0} \\ &= \left(\left(\prod_{C \in \text{mcl}(G)} \mu_{(\beta_0|_{C_0}) \| (\beta'|_{C'})}^C \right)_{\beta' \in \prod_{j \in (N_1 \cup N_2) \setminus N_0} [d_j]} \right)_{\beta_0} = \varphi_G(\mu), \end{aligned}$$

where, for $C \in \text{mcl}(G)$ and $\alpha \in \prod_{j \in C} [d_j]$, the parameter μ_α^C is defined as

$$\mu_\alpha^C := \begin{cases} \theta_\alpha^C & \text{if } C \subseteq N_1 \text{ and } C \not\subseteq N_0, \\ \eta_\alpha^C & \text{if } C \subseteq N_2 \text{ and } C \not\subseteq N_0, \text{ and} \\ \theta_\alpha^C \eta_\alpha^C & \text{if } C \subseteq N_0. \end{cases}$$

This computation proves that $X_{G_1} * X_{G_2} \subseteq X_G$. Conversely, given any parameter vector μ for G , we can let θ be the restriction of μ to maximal cliques of the first and third type above, and set η_α^C equal to μ_α^C if C is of the second type above and equal to 1 if it is of the third type. This yields the opposite inclusion. \square

Proof of Theorem 1. By Proposition 21, the ideal I_G is the ideal of the iterated toric fibre product $X_{G_1}^{*a_1} * \cdots * X_{G_s}^{*a_s}$. By Lemma 20, each of the varieties X_{G_i} is Hadamard closed. Hence Theorem 5 applies, and I_G is generated by polynomials of degree less than some D , which is independent of a_1, \dots, a_s . Then it is also generated by the binomials of at most degree D . \square

REFERENCES

- [DK14] Jan Draisma and Jochen Kuttler. Bounded-rank tensors are defined in bounded degree. *Duke Math. J.*, 163(1):35–63, 2014.
- [Dra10] Jan Draisma. Finiteness for the k-factor model and chirality varieties. *Adv. Math.*, 223:243–256, 2010.
- [DS98] Persi Diaconis and Bernd Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Ann. Stat.*, 26(1):363–397, 1998.
- [EKS14] Alexander Engström, Thomas Kahle, and Seth Sullivant. Multigraded commutative algebra of graph decompositions. *J. Algebr. Comb.*, 39(2):335–372, 2014.
- [HC71] J.M. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. Unpublished, <http://www.statslab.cam.ac.uk/~grg/books/hammfest/hamm-cliff.pdf>, 1971.
- [HS12] Christopher J. Hillar and Seth Sullivant. Finite Gröbner bases in infinite dimensional polynomial rings and applications. *Adv. Math.*, 221:1–25, 2012.
- [KR14] Thomas Kahle and Johannes Rauh. Toric fiber products versus Segre products. *Abh. Math. Semin. Univ. Hamb.*, 84(2):187–201, 2014.
- [Lau98] Steffen L. Lauritzen. *Graphical models*, volume 17 of *Oxford Statistical Science Series*. Oxford Univ. Press., Oxford, 1998.
- [Mac01] Diane MacLagan. Antichains of monomial ideals are finite. *Proc. Am. Math. Soc.*, 129(6):1609–1615, 2001.
- [Oos16] Florian M. Oosterhof. *Stabilisation of iterated toric fibre products*. 2016. Master’s thesis, Eindhoven University of Technology, <http://repository.tue.nl/855107>.
- [RS14] Johannes Rauh and Seth Sullivant. The Markov basis of $k_{3,N}$. 2014. Preprint, arXiv:1406.5936.
- [RS16] Johannes Rauh and Seth Sullivant. Lifting Markov bases and higher codimension toric fiber products. *J. Symb. Comp.*, 74:276–307, 2016.
- [Sam15] Steven V. Sam. Ideals of bounded rank symmetric tensors are generated in bounded degree. *Invent. Math.*, 2015. To appear, arXiv:1510.04904.
- [Sam16] Steven V. Sam. Syzygies of bounded rank symmetric tensors are generated in bounded degree. 2016. Preprint, arXiv:1608.01722.
- [Shi12] Takafumi Shibuta. Gröbner bases of contraction ideals. *J. Algebr. Comb.*, 36(1):1–19, 2012.
- [SS14] Steven V. Sam and Andrew Snowden. Gröbner methods for representations of combinatorial categories. *J. Am. Math. Soc.*, 2014. To appear, arXiv:1409.1670.
- [Sul07] Seth Sullivant. Toric fiber products. *J. Algebra*, 316(2):560–577, 2007.

(Jan Draisma) MATHEMATICAL INSTITUTE, UNIVERSITY OF BERN, SIDLERSTRASSE 5, 3012 BERN, SWITZERLAND
E-mail address: jan.draisma@math.unibe.ch

(Florian M. Oosterhof) DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, EINDHOVEN UNIVERSITY OF TECHNOLOGY, P.O. BOX 513, 5600 MB EINDHOVEN, THE NETHERLANDS

E-mail address: f.m.oosterhof@tue.nl